WHITEPAPER

IDENTIFYING SINUS RELATED MEDICAL CONDITIONS WITH AI-BASED DIAGNOSTIC TOOLS

Clinical decision support tools from NVIDIA, Microsoft Azure, Inform AI and SFL Scientific

Health problems are fundamental issues for people worldwide, so improving services, delivery methods, and diagnostic accuracy is critical for improving health outcomes. The collection and automation of health data plays an equally critical role in addressing health problems as the industry works towards automated diagnosis and support systems for better patient care.

Deep learning techniques, such as machine learning and artificial intelligence algorithms have a tremendous potential to influence the practice of radiology and medical imaging at large. Unlike many other fields of medicine, nearly all of the primary data collected in medical imaging is now digital, and with advances in modern electronic health records, cloud networking, infrastructure, and clinical data distribution, these digital data lend themselves to analysis by artificial intelligence algorithms. The result is an increasingly large amount of image data available to answer clinically meaningful questions and to develop advanced medical applications. Deep learning models can be used for detection, diagnosis, segmentation, and simulation using MRI, CT, X-ray, and digital pathology data. This boom in innovation is partly due to the significant advances made in computational processing of very large datasets. NVIDIA is at the forefront of this computing renaissance, with its Volta GPU architecture. Collectively, these advances constitute a very disruptive force that will create new markets and drive transformation for many healthcare organizations.

There have been several recent successes in applying deep learning to medical imaging tasks, notably in projects developing algorithms for the detection of diabetic retinopathy, detection of skin lesions using dermatoscopic images, detection of cancer metastases using pathology slides, and detection of numerous anatomical and bone fracture conditions in X-ray images. Medical imaging is an optimal opportunity for deep learning and algorithmic success because there is often a direct mapping of the input image pixel data to a specific diagnosis and increasing regulatory approval for AI-based tools. InformAI and SFL Scientific have partnered with the goal of accelerating the development of these novel tools by combining medical image analysis and deep learning. Together, the companies are focused on applying data-centric strategies and technology solutions to create clinical decision support tools with medical-grade accuracy that meet the confidence and consistency standards necessary for rigorous healthcare use.

These tools will assist clinicians with "augmented intelligence" by providing speed, filtering, and diagnostic capabilities, as well as helping mitigate the black-box mysticism of AI. InformAI and SFL Scientific are fortunate to be working on applications that have clinical impact and that improve outcomes. The two companies are poised at the tip of the iceberg of a market for scientists, physicians, and researchers to extract critical care information from increasingly available unstructured data.

Medical Applications for Deep Learning

Every year, nearly one million surgical procedures are performed to remedy sinusitis, and $10B is spent in the United States annually on medical care for sinusitis conditions. Surgeries are performed to address serious medical paranasal sinus conditions involving soft tissue inflammation, infections, nasal deviations, polyps, and tissue masses. Together with InformAI's R&D team, SFL Scientific has worked towards developing solutions to classify and predict sinus-related medical conditions and disease states from 3D computed tomography (CT) head scans. The goal is to create a suite of diagnostic tools centered around deep learning image classifiers to predict the occurrence of diseases and to assist radiologists and physicians in speeding up the process of identifying sinus related medical conditions.

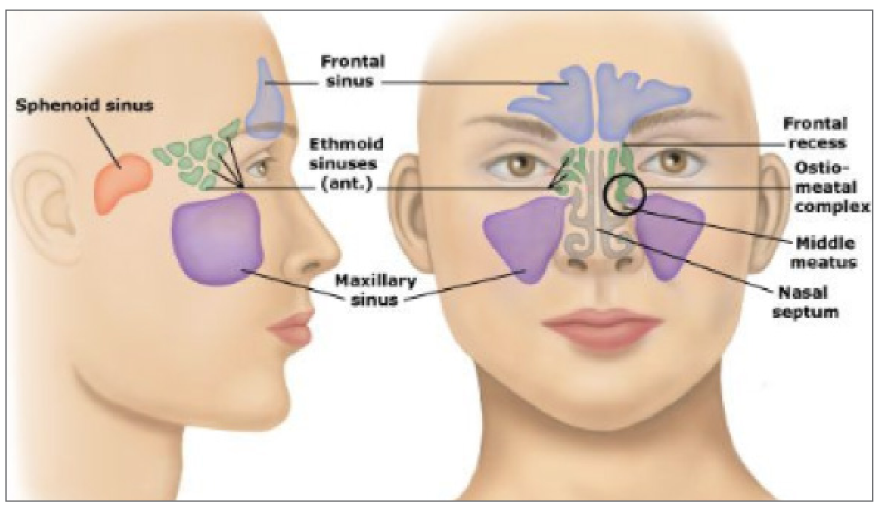

Figure-1

Deep learning applications often rely on extremely large datasets, however, the availability of ground truth data annotated by expert radiologists and pathologists is not easily available for many reasons, including privacy concerns, commercialization aspects, time and manpower considerations and, of course, cost. Simply put, annotation of medical data is expensive and tedious. For certain rare medical conditions, it is nearly impossible to collect enough representative data examples to use for training datasets. In late 2017, InformAI and SFL Scientific started working together to develop a deep learning-based system for computer-aided diagnostics, bringing a new level of intelligence to legacy imaging solutions and clinics. InformAI, a healthcare analytics company, builds business analytics and artificial intelligence-based solutions for industry-leading healthcare organizations, medical research institutions, and medical device companies to improve operational efficiency, patient outcomes, and medical diagnosis. InformAI, working with its clients, is focused on building a platform of medical image classifiers, surgical risk predictors, and patient outcome predictors. To accomplish these goals, InformAI has assembled some of the largest patient image datasets in the industry for AI model training development. InformAI is part of the NVIDIA Inception Program and a technology partner with Microsoft Azure. Together, InformAI and SFL Scientific, a data science consulting firm, are building a deep learning model to address medical conditions that affect the paranasal sinuses. SFL Scientific's capabilities include developing and implementing comprehensive algorithms, mining massive datasets, and deploying solutions at scale. The company provides integrated solutions and professional services that take advantage of NVIDIA's GPU computing and Microsoft Azure's platform portfolio. SFL Scientific is a preferred service provider for NVIDIA and choose Azure as the optimal environment to prototype deep learning algorithms, allowing R&D teams to accelerate adoption, work seamlessly across businesses, and deploy developed tools at scale. "The ability to couple large medical imaging datasets from our leading healthcare partners with the exceptional computational performance of the NVIDIA V100 GPUs on Azure Cloud, and novel deep learning model optimization from SFL Scientific have been critical in building our portfolio of AI image classifier applications," Jim Havelka, CEO, InformAI, explained.

Classification of Disease in 3D Medical Images

One of the biggest issues facing healthcare radiology professionals is the operational fatigue that comes with information overload and visual strain associated with the review of medical images. In fact, this problem represents the “weak link of a medical diagnosis that often slows down time to decisions. Furthermore, the physician fatigue may affect both the volume and quality of those patient decisions. To complicate matters, many parts of the world lack the volume of trained clinical staff with the expertise to diagnose these medical issues within their local communities. For all these reasons, computer-aided diagnostic tools can affect the overall success of patient treatment and potentially reduce medical errors by working as an unbiased “assist tool” to reduce variability and increase specificity of their readings. The amount of information available to clinicians is overwhelming, so reducing the time to identify abnormalities helps improve the productivity of clinicians and the associated diagnostic accuracy. Similarly, automating routine image examination stands to reduce the burden placed on clinical staff who are already pressed for time, especially in emergency or surgical scenarios.

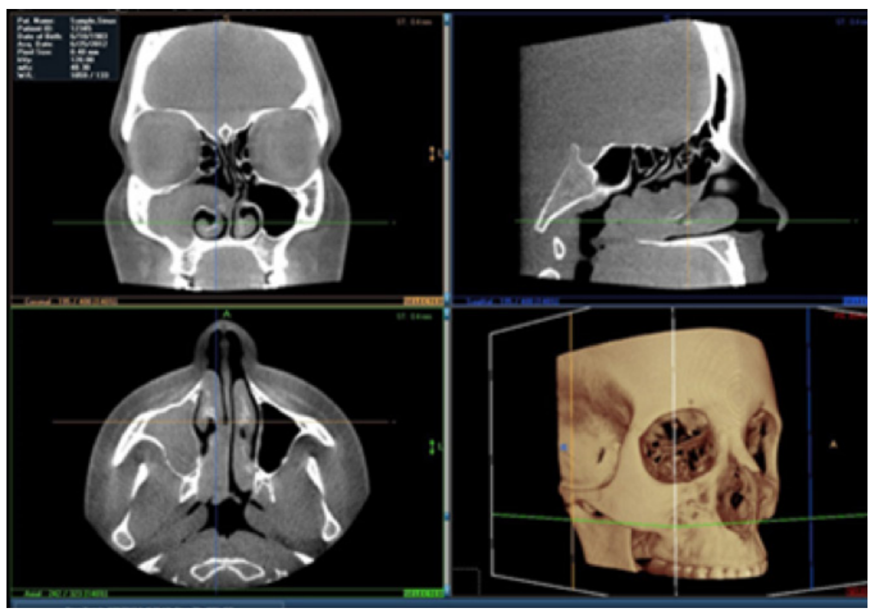

Figure-2

For machine learning, good predictive performance is only as good as the underlying criteria, data quality, and definitions of the annotations. InformAI created a framework to review and segment over 23 diseases within the sinus scans into specific regions and groups, greatly reducing the region of interest of the resultant 3D image stack. The non-scalability of maintaining an expert team, and the volume of data required to do so, motivated the development of learning methods that can use weakly-labeled training sets with global and binary labels in those regions to create algorithms for rapid differentiation of abnormalities from the healthy state. By feeding increasingly complex data sets, preprocessing, subsampling, and using augmentation techniques to computationally produce a larger volume of data, over time these algorithms begin to behave like a trained radiologist to quickly identify anomalies and provide confidence intervals for areas which will require human intervention from a clinician. Despite these recent advances, standalone tools having no physician involvement using direct applications of machine learning to healthcare remains a challenge inherent in the goal of making personalized predictions from large amounts of noisy, biased, and unstructured data.

Data Strategy and Deep Learning: Overcoming the Complexity of 3D Medical Imaging Data

Instead of the normal shades of gray that a human radiologist sees in CT scans, a computer represents each image as a matrix of numbers representing the pixel brightness. Traditional computer vision techniques typically involve computing the presence of numerical patterns in this matrix, such as boundaries for low-level features, and applying machine learning algorithms designed to distinguish images based on these features. Significant expertise and time are required to engineer the best features for distinguishing specific conditions and separating classes of images. This optimization problem of distinguishing features is traditionally difficult, but is the basis of deep learning; it uses hierarchical abstractions and different functional layers to learn representations and features from the data.

Convolutional neural networks (CNN) have been very successfully used for image classification and other types of imaging tasks. Convolution operations that extract image features produce matrices, and these "layers" in a CNN generate output matrices stacked in a volume. This volume can then serve as the input for another layer that may detect more complex features in the input image. Each layer can then be calculated and converted within the output nodes of the network into probabilities for classification. As no pre-trained deep networks are readily available for 3D image datasets for CT or MRI for baseline testing, CNNs for these images need to be trained from raw patient data with large labeled data sets to achieve the training/predict process shown in the figure below.

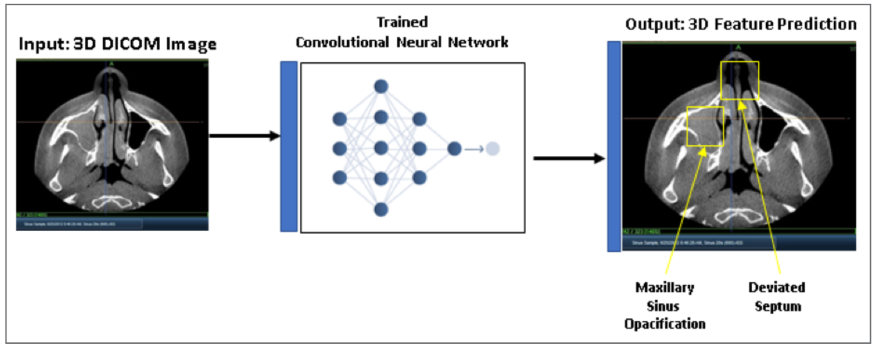

Figure-3

InformAI compiled large volumes of 3D CT images in DICOM image format from its partner network and recruited expert radiologists to assist in annotation and diagnosis, building the input training data set with extensive work to tune the locations of the disease zones in response to the feedback from the radiologist teams and clinical researchers. To begin the process of creating a deep learning pipeline, a suite of tools was developed to take patient anonymized studies and annotate and process the raw DICOM data into a suitable format for AI model development. Tools were developed to automate extracting the target 3D image segments from the series of 3D CT scans to prepare and ingest into the CNN Tensorflow models. The extracted segments correspond to regions that radiologists use to make diagnoses from different coronal, sagittal, and axial views of the human head. The software tools also reduced scan size, segmented the disease region of interest, lowered image noise to improve AI model prediction accuracy, and lowered the memory requirements and time required for training. As these voxel sizes are quite large, with scans approximately 400x400x300 pixels, the memory required surpassed what is typical available with consumer GPU hardware.

As part of this image preprocessing, registration algorithms were developed to accurately place anatomical features within regions of interest, resampling and normalizing the data in the process. As the disease features are very small, sometimes only a few voxels, training a CNN to detect and begin to generalize a particular disease in a volume of approximately 300,000 times the disease size requires many training epochs and computational scalability - using a typical computational platform to train these large CNN's is not practical. SFL Scientific partnered with NVIDIA and Azure teams on securing NC24s_v3 (24CPU/448 GB memory) instances, allowing for a network of NVIDIA V100 GPUs to serve as the computational platform that enabled project development. With the data ingestion and transfer requirements for this large volume of data, teams leveraged Microsoft Azure Blob Cloud Storage as a central repository, as using other storage schemes was not practical for maintaining the high data throughput required during training operations. By creating the code in Python and Tensorflow, the systems support transferability, ease of deployment, and operate in easily provisioned environments.

Together, InformAI and SFL Scientific developed the Deep Learning technology stack, involving 3D CNN models consisting of approximately 400 million parameters, to detect the targeted list of medical conditions as shown in the figure below. As accuracy of the final model is reliant on the quality of the dataset used for training the deep neural network, the extended curation, assembly, and pre-processing development was critical when considering performance.

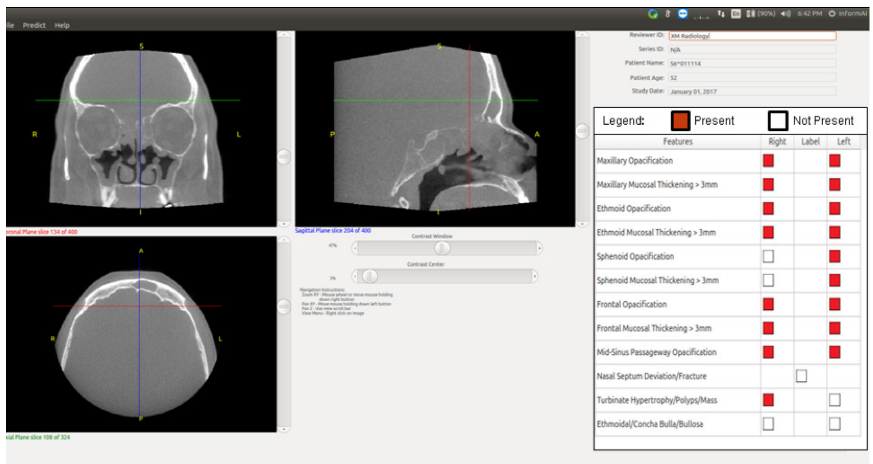

Figure-4

To refine the InformAI model, SFL Scientific conducted an extensive exploratory data analysis (EDA) on the datasets to understand the incidence and correlations of the medical conditions, concluding that despite the availability of thousands of annotated scans available, certain class imbalances, rare conditions, and positive examples of conditions of interest would still be underrepresented and would prevent achieving the necessary accuracies for viability in a real-world clinical setting.

To solve this problem, the companies used data augmentation techniques to increase the effective size of the training sets by exploiting all the symmetries of 3D space such as using random rotations, random transformations, random zooming by small amounts, mirroring, etc. Typically, there are two types of data augmentation that can be deployed: lossless and lossy. Given the 3D nature of the data, the lossless augmentation refers to the symmetries of the cube where rotations by 90º in any direction, brightness/contrast adjustments, or adding denoising steps preserves the fidelity of the image data. Lossy augmentation consists of small rotations (less than 90º), zooming, and down-sampling of that data for multiple views. While lossless augmentation can be generated quickly by exploiting the processing power of GPUs, as the data is being read into the network, lossy augmentation requires slow operations that create new data on storage as each 3D transformation needs to be individually computed and applied to each voxel.

For high-throughput performance, lossy data needs to be generated ahead of time and saved independently, effectively increasing the amount of data to be stored. By creating these various data assets with smaller input and features, curriculum learning can be employed to progressively train the model in the same manner that humans learn increasingly complicated tasks by incrementally adding difficulty to the learning. Specifically, learning performance is typically much better when the examples are not randomly presented in training, but organized in a meaningful order which illustrates gradually more complex concepts and produces a much more robust training strategy. Further, curriculum learning can be seen as a general strategy for global optimization of non-convex functions.

InformAI's Mission for Building Solutions

InformAI is focused on building AI-enabled tools to assist healthcare organizations in improving operational efficiency, patient outcomes and medical diagnosis. The company's products include complex 3D image classifiers for targeted medical conditions, surgical risk predictors, patient outcome predictors and operational decision-support tools. Our partners, SFL Scientific and NVIDIA, share that vision and provide AI model optimization and GPU computing capabilities that are instrumental to building these products. Further, with Microsoft Azure's flexible and enterprise-grade compute platform, we were able to assemble, ingest, and build highly accurate production tools with ease in a replicable environment. As we scale our products and services and integrate with our clients and partners, having the ability to manage solutions in real-time, across infrastructures, provides us the end to end framework we require.

NVIDIA's Impact on Medical Research & Development

Tesla V100 is engineered to provide maximum performance in existing hyperscale server racks. With AI at its core, the Tesla V100 GPU delivers 47X higher inference performance than a CPU server. This giant leap in throughput and efficiency makes the scale-out of AI services practical. NVIDIA's Volta GPU architecture addresses the great challenge of medical imaging: processing the massive sea of data.

"We're at a very fortunate time in medical imaging in that with all the technology NVIDIA has been developing over the past five years, we can now literally virtualize medical imaging pipelines."

Kimberly Powell, VP Healthcare NVIDIA

About NVIDIA:

NVIDIA's invention of the GPU in 1999 sparked the growth of PC gaming market, redefined modern computer graphics, and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI – the next era of computing – with the GPU acting as the brain of computers, robots, and self-driving cars that can perceive and understand the world.

About Microsoft Azure:

Microsoft (Nasdaq "MSFT" @microsoft) is the leading platform and productivity company for the mobile-first, cloud-first world, and its mission is to empower every person and every organization on the planet to achieve more.

About InformAI:

InformAI is a US-based healthcare analytics company that builds business analytics and artificial intelligence-based solutions for industry leading healthcare organizations, research institutions, and medical device companies to improve patient outcomes, operational efficiency and surgical risk predictors. InformAI, working with its clients, is focused on building a platform of medical image classifiers, surgical/patient risk predictors and operational decision-support applications. InformAI is part of the NVIDIA Inception Program.

About SFL Scientific:

SFL Scientific is a data science consulting and professional services company, providing a broad range of solutions in data engineering, machine learning, and Artificial Intelligence. We design, prototype, integrate, and manage sophisticated AI solutions by leveraging emerging technology. SFL Scientific creates opportunities for healthcare and life sciences organizations to deliver on data-driven diagnostic tools, products, patient care, and operational outcomes. SFL Scientific is a preferred deep learning service provider for NVIDIA and Microsoft Azure.